

CLAIMS

What is claimed is:

1. A computerized method of automatically generating from a first speech recognizer a second speech recognizer, said first speech recognizer comprising a first acoustic model with a first decision network and corresponding first phonetic contexts, and said second speech recognizer being adapted to a specific domain, said method comprising:
 - based on said first acoustic model, generating a second acoustic model with a second decision network and corresponding second phonetic contexts for said second speech recognizer by re-estimating said first decision network and said corresponding first phonetic contexts based on domain-specific training data.
2. The method of claim 1, wherein said domain-specific training data is of a limited amount only.
3. The method of claim 1, said re-estimating comprising:
 - partitioning said training data using said first decision network of said first speech recognizer.
4. The method of claim 3, said partitioning step comprising:
 - passing feature vectors of said training data through said first decision network and extracting and classifying phonetic contexts of said training data.
5. The method of claim 4, said re-estimating further comprising:
 - detecting domain-specific phonetic contexts by executing a split-and-merge methodology based on said partitioned training data for re-estimating said first decision network and said first phonetic contexts.

EXPRESS MAILING LABEL NO. EK972214875US

- 1 6. The method of claim 5, wherein control parameters of said split-and-merge
2 methodology are chosen specific to said domain.
- 1 7. The method of claim 5, wherein for Hidden-Markov-Models (HMMs) associated
2 with leaf nodes of said second decision network, said re-estimating comprises re-
3 adjusting HMM parameters corresponding to said HMMs.
- 1 8. The method of claim 7, wherein said HMMs comprise a set of states s_i and a set
2 of probability-density-functions (PDFS) assembling output probabilities for an
3 observation of a speech frame in said states s_i , and wherein said re-adjusting step is
4 preceded by:
5 selecting from said states s_i a subset of states being distinctive of said domain;
6 and
7 selecting from said set of PDFS a subset of PDFS being distinctive of said
8 domain.
- 1 9. The method of claim 7, wherein said method is executed iteratively for additional
2 training data.
- 1 10. The method of claim 8, wherein said method is executed iteratively for additional
2 training data.
- 1 11. The method of claim 7, wherein said first and said second speech recognizer are
2 general purpose speech recognizers.
- 1 12. The method of claim 7, wherein said first and said second speech recognizers

1 are speaker-dependent speech recognizers and said training data is additional speaker-
2 dependent training data.

1 13. The method of claim 7, wherein said first speech recognizer is a speech
2 recognizer of at least a first language and said domain specific training data relates to a
3 second language and said second speech recognizer is a multi-lingual speech
4 recognizer of said second language and said at least first language.

1 14. The method of claim 1, wherein said domain is selected from the group
2 consisting of a language, a set of languages, a dialect, a task area, and a set of task
3 areas.

1 15. A machine-readable storage, having stored thereon a computer program having
2 a plurality of code sections executable by a machine for causing the machine to
3 automatically generate from a first speech recognizer a second speech recognizer, said
4 first speech recognizer comprising a first acoustic model with a first decision network
5 and corresponding first phonetic contexts, and said second speech recognizer being
6 adapted to a specific domain, said machine-readable storage causing the machine to
7 perform the steps of:

8 based on said first acoustic model, generating a second acoustic model with a
9 second decision network and corresponding second phonetic contexts for said second
10 speech recognizer by re-estimating said first decision network and said corresponding
11 first phonetic contexts based on domain-specific training data.

1 16. The machine-readable storage of claim 15, wherein said domain-specific training
2 data is of a limited amount only.

1 17. The machine-readable storage of claim 15, said re-estimating comprising:

1 partitioning said training data using said first decision network of said first speech
2 recognizer.

1 18. The machine-readable storage of claim 17, said partitioning step comprising:
2 passing feature vectors of said training data through said first decision network
3 and extracting and classifying phonetic contexts of said training data.

1 19. The machine-readable storage of claim 18, said re-estimating further comprising:
2 detecting domain-specific phonetic contexts by executing a split-and-merge
3 methodology based on said partitioned training data for re-estimating said first decision
4 network and said first phonetic contexts.

1 20. The machine-readable storage of claim 19, wherein control parameters of said
2 split-and-merge methodology are chosen specific to said domain.

1 21. The machine-readable storage of claim 19, wherein for Hidden-Markov-Models
2 (HMMs) associated with leaf nodes of said second decision network, said re-estimating
3 comprises re-adjusting HMM parameters corresponding to said HMMs.

1 22. The machine-readable storage of claim 21, wherein said HMMs comprise a set
2 of states s_i and a set of probability-density-functions (PDFS) assembling output
3 probabilities for an observation of a speech frame in said states s_i , and wherein said re-
4 adjusting step is preceded by:

5 selecting from said states s_i a subset of states being distinctive of said domain;

6 and

7 selecting from said set of PDFS a subset of PDFS being distinctive of said
8 domain.

1 23. The machine-readable storage of claim 21, wherein said method is executed
2 iteratively for additional training data.

1 24. The machine-readable storage of claim 22, wherein said method is executed
2 iteratively for additional training data.

1 25. The machine-readable storage of claim 21, wherein said first and said second
2 speech recognizer are general purpose speech recognizers.

1 26. The machine-readable storage of claim 21, wherein said first and said second
2 speech recognizers are speaker-dependent speech recognizers and said training data
3 is additional speaker-dependent training data.

1 27. The machine-readable storage of claim 21, wherein said first speech recognizer
2 is a speech recognizer of at least a first language and said domain specific training data
3 relates to a second language and said second speech recognizer is a multi-lingual
4 speech recognizer of said second language and said at least first language.

1 28. The machine-readable storage of claim 15, wherein said domain is selected from
2 the group consisting of a language, a set of languages, a dialect, a task area, and a set
3 of task areas.